Proceedings of the IEEE ICRA 2009
Workshop on People Detection and Tracking
Kobe, Japan, May 2009

Experimental Evaluation of a People Detection Algorithm in Dynamic Environments

Dario Lodi Rizzini, Stefano Caselli
RIMLab - Robotics and Intelligent Machines Laboratory
Dipartimento di Ingegneria dell'Informazione
University of Parma, Italy
E-mail {dlr,caselli}@ce.unipr.it

***Abstract*— People detection is an important capability both for human-robot interaction in service robotics and to distinguish the stable environment from the perturbation due to people motion in localization and mapping tasks. Several techniques have been proposed for different application contexts and sensors. Range data acquired by laser scanners are metrically accurate and suitable for computationally-inexpensive people detection. Furthermore, laser scans provide a geometric description of local environment that can be combined with the information about dynamic objects.**

In this paper, a previously proposed method for detecting people legs from laser scans is experimentally evaluated and exploited to improve scan matching by removing dynamic parts corresponding to people. This algorithm splits laser scans into beam segments and classifies each segment. Classifications of simple features are then combined into a boosted classifier with Adaboost. The fundamental assumption of scan matching is that consecutive scans can be aligned with a rigid body transformation, since they represent the same scene. When dynamic elements like human legs are removed from scans, such assumption holds. We also investigate the effectiveness of the proposed people detection algorithm in terms of its ability to generalize across different environments and to support track persistency across scans.

I. INTRODUCTION

The aim of service robotics is the execution of tasks for people care. As a mobile service robot moves in an environment populated by people, robot-human interaction is therefore a fundamental requirement. Furthermore, even if the tasks to be performed do not involve people care, the recognition of dynamic elements including people is required for localization and mapping. Indeed, localization and mapping algorithms usually assume the complete state hypothesis. According to this hypothesis, the evolution of the system consisting of the robot and the static environment is completely described by the state variables. State usually includes robot location and map descriptors, but it does not consider human presence. Thus, such assumption is strongly violated in populated environments. Solutions for this problem range from filtering the dynamic obstacles to classifying and tracking them.

Several approaches have been proposed for people detection depending on the available sensor data and the context of application. The most popular sensors are cameras and range finders. Range finders have the advantage of limited processing requirements. Limiting our survey to laser-based robot applications, the approaches can be divided into tracking oriented techniques and geometric rule classifiers. The first category includes simple extensions of localization or SLAM algorithms [1], [2], [3] or specifically designed techniques [4], [5]. The second category includes all the methods that perform a classification using the features extracted from laser scans [6], [7], [8]. However, the above categorization remains arbitrary since tracking and feature detection are typically mixed together.

In this paper, we experimentally evaluate the algorithm for detecting people proposed in [7] that combines several feature based classifiers to perform a more robust estimation according to Adaboost boosting technique. This method has the advantage of performing people detection on a single scan without depending on a specific tracking technique or on assumptions about motion of people.

Furthermore, we use this algorithm to improve scan matching performance in a populated environment and apply the concept of track persistency to the classification results. The fundamental assumption of scan matching is that consecutive scans can be aligned with a rigid body transformation, since they represent the same scene. As discussed before, this assumption is violated in a populated environment, but the people detection algorithm can be used to filter out people presence. Our contribution lies in the experimental evaluation of the robustness of a scan matching technique and of the improvement allowed by people filtering. A further application of people classification relies on the concept of persistency [8]. A track corresponding to a person is persistent if the segment associated to the given track in each scan is often classified correctly. The evaluation of the persistency of people tracks yielding the potential of speeding-up the training of the classifier is the final contribution of the paper.

The paper is organized as follows. Section II briefly describes the algorithm for people detection. Section III illustrates the application of the classifier to improve the scan matching problem and the possibility to exploit track persistency for semi-unsupervised training. Section IV presents the experimental results. Finally, section V summarises the paper drawing some conclusions and perspectives.

II. PEOPLE DETECTOR

This section illustrates the algorithm for detecting people and its application to recognize dynamic and stable elements in the environment. The basic people detection algorithm has been adapted from [7], as described next. The algorithm operates on a single laser scan in order to find if any subset in the range readings of the scan corresponds to a person as described in the following. First, the scan is divided into groups of adjacent range values called segments. Second, the algorithm classifies the segments establishing their correspondence to people legs. The classifier is achieved by combining several elementary classifiers that operate by extracting a specific feature from the segment and evaluating the value of such feature.

In literature, the outlined method for combining weak decisors in order to reduce the classification error is known as *boosting*. Adaboost algorithm is one of the most extensively used boosting algorithm [9]. The input of the algorithm is the training set, a set of examples (the scan segments in this case) labeled with the result of correct classification. Adaboost builds the final classifier by iteratively executing a learning round. During each round, the weak classifiers are trained using the examples of training set and the classifier that minimizes the classification error is selected for the round. The classification error is computed by weighting the error of each misclassified example. Weights are larger for the examples that have been wrongly classified in previous rounds. The classification error is then used to compute the coefficient that measures the contribution of the weak classifier to the decision.

Adaboost is a meta-algorithm that does not impose the form of the weak classifier. For people detection based on laser scans, since the features extracted from each segment are represented by a scalar, the weak classifiers $h_j(\cdot)$ have the following fixed expression

$$h_j(e) = \begin{cases} true & \text{if } \mathrm{p_j f_j(e)} < \mathrm{p_j}\theta_\mathrm{j} \\ false & \text{otherwise} \end{cases} \tag{1}$$

where e is the item to be classified (the segment), $f_j(e)$ is the feature extracted from e, θ_j is the decision threshold and $p_j \in \{+1, -1\}$ gives the direction of inequality. This form is suggested in [10] and adopted for the people detector in [7].

A. Feature Definition

The features used in the described classifier are scalar values computed from a scan segment. As explained above, a segment is a set of consecutive range values of a laser scan approximately corresponding to a distinguishable object of the environment. Segmentation is an important step of the algorithm, which is sometimes neglected. In the experiments section, it will become apparent how segmentation affects the final result. In this paper, a simple splitting technique has been used. The range values of the scan are traversed in counterclockwise order and, when the jump distance of a range reading with respect to the previous reading is above a threshold, a new segment starts. Segments including only one range reading are discarted. Comparing with the original proposal in [7], it is unclear whether our segmentation technique exactly reproduces the original approach; if not, this is the only significant difference between the original algorithm and our implementation.

The range values of the segment are then transformed into cartesian coordinates with respect to the local reference frame fixed on the sensor. Depending on the feature, polar or cartesian coordinates are used. For each segment, we used the same set of 14 features proposed in the original paper, that are listed in the following.

1) *Number of points*.
2) *Standard deviation*: it is the mean distance from the mean value of the points of the segment.
3) *Mean average deviation from median*: it is a more robust version of the previous feature that uses the median point instead of the mean point. The median point coordinates are given by the 0.5 percentiles of the distribution of x and y coordinates of points.
4) *Jump distance from preceding segment*.
5) *Jump distance to succeeding segment*.
6) *Width*: it is the Euclidean distance between the first and the last point.
7) *Linearity*: it is the sum of square distances between each segment point and the regression line computed using the same points.
8) *Circularity*: it is the sum of square distances between each segment point and the regression circle computed using the same points. The regression circle is achieved according to least square criterion. When only two points are available, circularity is set to a large value.
9) *Radius*: it is the radius of the regression circle. When only two points are available, the radius is set to a large value.
10) *Boundary length*: it is the sum of the distances between consecutive segment points. It corresponds to the length of the boundary defined by the poli-line that connects each pair of points.
11) *Boundary regularity*: it is the standard deviation of the line
12) *Mean curvature*: it is the average value of the curvatures computed on triplets of consecutive points.
13) *Mean angular difference*: it is the average value of the angles computed on triplets of consecutive points.
14) *Mean speed*: it is the average speed of the range readings of the segment. The computation of range speed requires the value of the given range reading on the current and previous scans and the time interval between the acquisition of the two scans. Mean speed is the only feature that requires temporal correlation between two consecutive scans.

III. Multiple Scans Applications

The algorithm for detecting people described above has the remarkable advantage of performing a classification using only the geometric information available in a single scan, without requiring temporal correlations between scans. The only exception is represented by feature 14 that usually gives a negligible contribution to the boosted classifier as will be shown in section IV. Indeed, an algorithm that does not rely on temporal correlations is easier to implement and to test, since there is no specific constraint on the order of the scans. Moreover, the detection is independent from the motion state of the people and of the robot carrying the laser scanner. Common experience suggests that a person usually moves in an environment, but a robust people detector cannot rely on this assumption. In contrast to other techniques exploiting tracking, the illustrated algorithm does not require arbitrary dynamic models. However, this method can be easily integrated into a tracking system. In this subsection, we describe the application of the people detector to two different problems both related to temporal proximity: alignment of scan pairs and scan segment tracking.

A. People Filtering

Scan matching is the problem of finding a rigid motion that makes a laser scan overlap another reference scan. The fundamental assumption is that the two scans to be aligned share the representation of a region of the environment. Such hypothesis usually holds when the second scan is collected from a location near to the reference location and the environment is static. However, if there are people moving in the environment, such assumption is clearly violated. While several scan matching algorithms may be robust to such violations, the people detector can improve the performance of the scan matcher by removing the perturbation caused by human presence. We call this operation *people filtering* hence after. The effects of such correction are not easy to illustrate and to evaluate. First, if the motion of an object is too slow when compared to the frequency of acquisition, the object appears still in two consecutive scans. Second, the scan matcher can recover the values of translation and rotation from the fixed background that often dominates the scans. More details and results on people filtering are reported in section IV.

B. Track Persistency

A second application of the people detection algorithm exploits temporal correlation between scan segments to minimize the acquisition cost of training set. The described boosted classifier learns the value of internal parameters (the thresholds of weak classifiers, the weights, etc.) in a supervised training phase. Currently, the examples in the training set are manually labeled, but manual classification is a tedious and time-consuming operation. It would be convenient to perform a partially automatic labelling of the collected segments, at least to expand the existing training set. The concept of *track persistency* proposed in [8] could be used for this purpose. The original aim of this proposal is the unsupervised training of moving obstacles classifiers in a multi-sensor architecture. First, moving obstacles are detected as persistent tracks in the data acquired from a given sensor source. Second, these data are labeled as moving obstacles and are used to train a classifier.

Persistency can be applied to improve the performance of the described feature-based classifier learnt from an initial training set. In a typical scenario, one or more people move in a trajectory and their legs are repeatedly observed by a range finder. A perfect leg detector would find a segment for each person (or two segments, if the legs are distinguishable) in every scan and it would be possible to associate such segment to another segment corresponding to the same person in the previous scan. Thus, a persistent track could be found for each person and for fixed obstacles. Such correspondences between segments are not found in one of these cases:

- when the segmentation is not properly done;
- when the tracked person exits the visibility area;
- when the tracking algorithm fails;
- when the leg detector wrongly classifies the current or previous segment.

The latter case is the most interesting one because, if a wrongly classified segment is detected, it can be added to the training set and used to train a better classifier.

IV. Results

The aim of this section is to report the experimental evaluation of the legs detector described in previous sections and of the correction on scan matching error achieved with such algorithm. The experimental setup consists of an ActivMedia Pioneer I equipped with a Sick LMS 200 laser scanner. The scanning plane is $29.7\ cm$ high from the ground floor.

Experiments reported in the following have been performed in the Computer Engineer building of the University of Parma. The robot moved in different positions to collect scans from different locations in the environment during both training and evaluation steps. Figures 1(a)-(b) show two settings used in the tests: the hallway of Computer Engineering building and the Robotics laboratory. The two settings capture different kinds of rooms: the hallway is long and narrow and allows robot motions; the laboratory is full of obstacles, table and chairs legs that can be mistaken for human legs. The choice of the two rooms is similar to the one suggested in [7]. The main hallway of the Faculty building in Figure 1(c) was used only for the scan matching tests, since a larger environment was required. The robot moved to several places for each locations to collect training set data, but it stayed at a fixed position during the acquisition. The robot moved only during scan matching tests. Since classification is performed on a single scan, the motion of the robot is not significant for performance assessment.

The method described in this paper has been implemented independently from the original version described in [7]. The illustration given in the paper was sufficient to reimplement the same algorithm. Thus, the results shown in this section provide an independent validation for such technique. The differences between the two versions may depend only on the value of few parameters and on the segmentation procedure. A scan is split into a new segment when the jump between two consecutive ranges is greater than a given threshold, that has been set to $18\ cm$ for these experiments. Such a simple solution works quite well in almost all the considered cases, even if segments representing legs are sometimes confused with the background.

(a) (b) (c)

Fig. 1. Views of the experimental environments: (*a*) hallway of the Computer Engineering building; (*b*) laboratory; (*c*) hallway of the Faculty building.

A. Experiments with People Detector

The first set of experiments has been devoted to the assessment of leg detector performance. Scans have been acquired from the hallway and the laboratory in Figure 1(a)-(b) with people inside as discussed above. A third collection of scans has been acquired in the hallway after inserting additional obstacles of several sizes and shapes, since the hallway contained few obstacles during the first acquisition. Thus, three settings will be initially considered: hallway, hallway with obstacles and laboratory. During the experiments one or two people moved in the area.

The acquired training set and test set contain respectively 300 and 341 scans. The number of segments extracted from the training set is 5798, but only 2713 segments contain more than one point. In the global test set, there are 3315 segments consisting of more than one point on a total amount of 10259. This result is a consequence of the simple segmentation technique that splits when the jump distance is above the threshold. In order to improve the efficiency of the classifier and to avoid classification of segments with a single range reading, we considered as eligible segments only those with more than one range reading value.

	Detected Label (Hallway Training Set)		
True Label	Person	No Person	Total
Person	454 (90.98%)	45 (9.02%)	499
No Person	84 (8.76%)	875 (91.24%)	959
	Detected Label (All Training Set)		
True Label	Person	No Person	Total
Person	424 (84.97%)	75 (15.03%)	499
No Person	94 (9.80%)	865 (90.20%)	959

TABLE I

RESULTS OF PEOPLE DETECTION IN THE HALLWAY WITH FEW OBSTACLES.

Tables I, II and III show the results for the hallway, the hallway with more obstacles and the laboratory. In these three tables, the top part provides the results obtained with the classifier trained with the data of the specific environment and the bottom part the results obtained with the timing data collected from all the environments. In all three settings, the latter classifier generally performs worse than the specifically trained one, which performs correct detection, in average, on 90% of cases. The globally trained classifier only seems to

	Detected Label (Hallway Obs. Training Set)		
True Label	Person	No Person	Total
Person	93 (83.78%)	18 (16.22%)	111
No Person	24 (4.26%)	539 (95.74%)	563
	Detected Label (All Training Set)		
True Label	Person	No Person	Total
Person	107 (96.40%)	4 (3.60%)	111
No Person	263 (46.71%)	300 (53.29%)	563

TABLE II

RESULTS OF PEOPLE DETECTION IN THE HALLWAY WITH OBSTACLES.

	Detected Label (Laboratory Training Set)		
True Label	Person	No Person	Total
Person	143 (89.94%)	16 (10.06%)	159
No Person	135 (13.18%)	889 (86.82%)	1024
	Detected Label (All Training Set)		
True Label	Person	No Person	Total
Person	146 (91.82%)	13 (8.18%)	159
No Person	277 (27.05%)	747 (72.95%)	1024

TABLE III

RESULTS OF PEOPLE DETECTION IN THE LABORATORY.

reduce the number of false negatives for the hallway with obstacles and the laboratory, but it increases the number of false positives. We remark that the statistics illustrated above do not include the one-point segments that are filtered before performing the classification. Otherwise, the number of correct "no people" classifications for hallway, hallway with obstacle and laboratory would increase respectively of 1638, 691 and 1021.

	Test Set		
Training Set	Hallway	Hallway Obs.	Laboratory
Hallway	91.15%	64.10%	72.44%
Hallway Obs.	83.20%	93.77%	73.37%
Laboratory	83.54%	68.99%	87.24%

TABLE IV

COMPARISON OF TRAINING SETS.

In order to gain some insight into the potential for environment generalization of the people detection algorithm, Table IV compares the percentage of correct classification achieved with classifiers learnt from different training sets. The hallway with obstacles and laboratory classifiers provide the best global performance, hinting that richer environments should be used to favor generalization. The features that

Environment	Five best features
Hallway	9, 4, 4, 3, 7
Hallway Obs.	9, 7, 3, 11, 13
Laboratory	4, 3, 12, 9, 7
All	2, 7, 9, 7, 3

TABLE V

THE BEST FIVE FEATURES FOR EACH CLASSIFIER.

allow better results (Table V) are *radius* (9), *mean average deviation from median* (3), *jump distance* (4), and *linearity* (7). They are almost the same features reported in [7]. From a general viewpoint, our experimental results confirm with independent implementation and assessment, the results reported in [7].

B. Evaluation of People Filtering

The aim of the second set of experiments is the evaluation of the impact of people dynamics on operations that assume a static world. In particular, the described leg detector can be directly exploited for all the methods that extract geometric information from laser scans. For example, scan matching allows the estimation of local robot motion by aligning a pair scans acquired in two different locations. Scan matching presumes that two consecutive scans overlap on the common region when the correct rigid motion is applied. However, if the two scans contain segments corresponding to dynamic objects like people, the relative position between these segments and the environment may change.

In these experiments, the illustrated classifier is used to filter the scan segments corresponding to legs that should not be considered in scan alignment. The robot moved with a mean speed of $0.2\ m/s$ acquiring a laser scan approximatively every $100\ ms$. Experiments were performed in two environments. The first environment is the hallway of the Faculty shown in Figure 1(c). Since the people leg detector was not trained in this setting, we used the classifier trained with the Hallway Obs. dataset. The size of this environment allowed the robot to cover a path of about $25\ m$. The second environment is the hallway of the Computer Engineering building (Figure 1(a)), where classifier performance was tested. One or two people were walking in the environment at moderate speed. A standard scan matcher based on *iterative closest point* (ICP) algorithm [11] has been used. Since no ground truth information was available, the final robot pose estimated using scan matching on filtered scans has been compared with the final pose estimated on raw scans. These data do not represent a real error, but a displacement between two different evaluations.

Table VI illustrate the displacements for each coordinate of robot pose obtained in the experiments in the Faculty hallway. The overall position displacement is $9.3\ cm$ on a distance of about $25\ m$. Thus, the scan matcher is only affected to a limited extent by the presence of people. A second experiment was performed in a setting where the

Final Coordinate	$x\ (m)$	$y\ (m)$	$\theta\ (m)$
People Filtering	16.026	9.006	0.0114
No People Filtering	15.942	8.967	0.0180

TABLE VI

RESULTS OF SCAN MATCHING WITH PEOPLE FILTERING IN THE FACULTY HALLWAY.

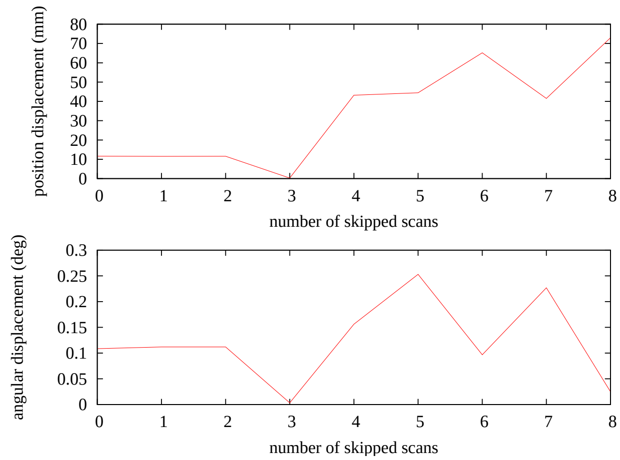

Fig. 2. Displacements between the final position (top) and between the final angle (bottom) of robot estimated by scan matching with leg detection enabled and disabled in Computer Engineering building. The displacement changes with the number of skipped scans.

people leg detector has been trained and tested. Figure 2 shows the position (top) and the angular (bottom) displacements varying with the number of skipped scans for a single experiment. In several cases, the human motion is slow when compared with the frequency of sensor acquisition and dynamic objects could be considered approximately static in two consecutive scans. Scans are then skipped to simulate systems subjected to computational load that cannot perform scan matching between all the pairs of scans or people moving at a faster rate. Note that the angular displacement is negligible even if the number of skipped scans is increased. Thus, orientation is not affected by people, at least in an environment like the considered hallway that has the strong reference provided by the parallel walls. On the other hand, the position displacement increases from about $1\ cm$ to $7\ cm$, when 8 scans are skipped, even though not monotonically. Thus, the evaluation of position is less robust than the evaluation of orientation, but scan matching is only marginally improved by people removal.

We interpret these results as follows. Motion estimation techniques like those based on scan matching may have the ability to filter out people presence, especially when only few slowly-moving people occlude a small portion of the scan. However, people filtering technique may play an important role for more cluttered and complex environments.

C. Towards Semi-unsupervised Labelling

The third set of experiments is devoted to the evaluation of track persistency as a criterion for semi-unsupervised segment labelling.

For this experiment, a simple tracking algorithm has been implemented. Each segment belonging to a scan is associated to the nearest segment of the previous scan. The considered distance is the distance between the centers of the two segments. Such naive association criterion is sufficient to achieve the results illustrated in the following, but a better segment matching would improve performance. For example, an accurate association rule should evaluate whether more segments correspond to the same object, e.g. two legs belonging to the same person. The robot moved in the hallway (Figure 1(a)) and acquired laser scans from the environment, while one or two people wandered in front of the range finder. Robot motion is estimated by matching pairs of consecutive scans and each scan is filtered removing the segments corresponding to legs as explained above. Such estimation is used to move the segments of the previous scan before performing the association.

Sequence number	Track Persistency Percentage People	Track Persistency Percentage No People
1	79.13%	96.96%
2	83.23%	96.87%
3	92.27%	96.82%
4	89.29%	97.14%

TABLE VII

TRACK PERSISTENCY PERCENTAGE

The first parameter evaluated is the persistency of people tracks and no people tracks. A segment with a given classification is called persistent if it is associated to another segment with the same classification. The persistency of a category track can then be measured by the ratio between the number of persistent segments and the total number of segments belonging to this category. Table VII shows the persistency ratio of people tracks and of no people tracks for four scan sequences acquired in the hallway. People persistency is about 80% for two sequences and 90% in the other two sequences. Such high values demonstrate that both the classifier and the tracking system work quite well.

However, we are interested in the remaining $10 - 20\%$ of positively classified segments that are associated to negatively classified segments in the previous scan. Such negative segments are possibly false negative. Currently, the tracking algorithm is not sufficiently accurate to make a decision and to add them to the training set.

V. CONCLUSION

In this paper, we experimentally evaluated an algorithm for detecting people based on boosted features and tested two multiple scan applications. In particular, we found that the people detector performs a correct classification in about 90% of cases, when the training set has been acquired in the same environment of the test set. The percentage decreases when a different training set is used. Differences between the results illustrated in this paper and in [7], where the algorithm was proposed, may be related to the segmentation method.

Furthermore, the classifier has been applied to filter people presence and improve scan matching in a populated environment. Experimental results demonstrate that scan matching is robust to the violation of the static environment assumption and that people filtering marginally modifies the estimation. However, further experiments in cluttered and complex environments are likely to emphasize the benefits provided by the people detector.

The third contribution of this paper is the experimental evaluation of track persistency over a sequence of scans. A track corresponding to a person is persistent if the segment associated to the given track in each scan is often classified correctly. Experiments illustrate that the people detection algorithm recognizes tracks with high persistency and only in few cases a person segment is not associated to another segment classified in the same way. Such track interruptions are due to several reasons, but may correspond to false positives of the algorithm. Therefore, persistency may be used to improve people detection by collecting misclassified segments for further training.

VI. ACKNOWLEDGMENTS

The authors gratefully thank Filippo Agazzi, Paolo Bertacchini and Gabriele Novelli for their help in setting up the experiments. This reaserch is partially supported by laboratory AER-TECH of Regione Emilia-Romagna, Italy.

REFERENCES

[1] C. Wang, C. Thorpe, and S. Thrun, "Online simultaneous localization and mapping with detection and tracking of moving objects: Theory and results from a ground vehicle in crowded urban areas," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2003.

[2] D. Haehnel, W. Burgard, and S. Thrun, "Map building with mobile robots in dynamic environments," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2003.

[3] D. Schulz, W. Burgard, D. Fox, and A. Cremers, "People tracking with mobile robots using sample-based joint probabilistic data association filters," *Journal of Robotics & Autonomous Systems*, vol. 12, no. 3, pp. 99–116, 2003.

[4] A. Fod, A. Howard, and M. Mataric, "Laser-based people tracking," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2002.

[5] K. Arras, S. Grzonka, M. Luber, and W. Burgard, "Efficient people tracking in laser range data using a multi-hypothesis leg-tracker with adaptive occlusion probabilities," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2008.

[6] J. Xavier, M. Pacheco, D. Castro, and A. Ruano, "Fast line, arc/circle and leg detection from laser scan data in a player driver," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2005.

[7] K. Arras, O. Martinez Mozos, and W. Burgard, "Using boosted features for the detection of people in 2D range data," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2007.

[8] R. Katz, B. Douillard, J. Nieto, and E. Nebot, "A self-supervised architecture for moving obstacles classification," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2008.

[9] R. Shapire and Y. Singer, "Improved boosting algorithms using confidence-rate predictions," *Machine Learning*, vol. 37, no. 3, pp. 297–336, 1999.

[10] P. Viola and M. Jones, "Robust real-time objects detection," in *Proc. of IEEE Workshop on Statistical Theories of Computer Vision*, 2001.

[11] A. Censi, "An ICP variant using a point-to-line metric," in *Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA)*, 2008.